RAND Journal of Economics
Vol. 32, No. 3, Autumn 2001
pp. 408–427

Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey

James H. Cardon*

and

Igal Hendel**

Adverse selection is perceived to be a major source of market failure in insurance markets. There is little empirical evidence on the extent of the problem. We estimate a structural model of health insurance and health care choices using data on single individuals from the NMES. A robust prediction of adverse-selection models is that riskier types buy more coverage and, on average, end up using more care. We test for unobservables linking health insurance status and health care consumption. We find no evidence of informational asymmetries.

1. Introduction

■ Since the seminal work by Akerlof (1970), extensive attention has been devoted to informational asymmetries among economic agents. The main consequence of such asymmetries is market failure. In spite of the extensive theoretical interest, there is little empirical evidence on the extent of the problem. The goal of this article is to test for adverse selection in the health insurance market.

Insurance markets are likely to suffer from adverse selection. Individuals may know more about their risk types than insurers (Rothschild and Stiglitz, 1976). Inefficiencies in health insurance markets are of significant policy concern, since part of the population may be uninsured. In the United States, a substantial fraction of the population has no health insurance coverage. Another observed symptom, consistent with the theoretical predictions, is that the uninsured tend to work for small employers. Large employers can overcome adverse selection by risk pooling. In spite of being the textbook example of a market with adverse selection, evidence on the importance of asymmetric information in insurance markets is inconclusive.[1]

* Brigham Young University; cardon@byu.edu.

** University of Wisconsin-Madison and NBER; igal@ssc.wisc.edu.

We would like to thank Pierre-André Chiappori, Penny Goldberg, Bo Honoré, Jim Powell, Steven Stern, Rob Porter (the Editor), two referees, and various seminar participants for useful comments. Financial support from the Department of Labor is gratefully acknowledged. The opinions expressed are the sole responsibility of the authors and do not represent the views of the U.S. Department of Labor.

[1] "There is as yet little direct evidence on whether or not adverse selection is or must be an important problem in health insurance, largely because it is difficult to define any kind of strong test" Pauly (1986, p. 650). For the literature on other insurance markets see Chiappori and Salanié (2000a).

Testing contract-theoretic models of asymmetric information is difficult. These models involve unobservable actions or types that are typically not observable to the researcher. The main components in these models are the contracts offered, the performance of the agents, and the transfers between the parties. One can think of testing the predictions of the theory in two different ways. One way is to ask whether the set of offered contracts conforms to the theoretical predictions. The second is to look at the agents' behavior given the set of contract options they face. The main hurdle to empirical work is the lack of appropriate data on contracts as well as on performance. Consider, for example, a labor relationship. Many elements of labor contracts are implicit, and therefore difficult to uncover in empirical work. Moreover, the performance of individual workers is difficult to measure, because the individual's output cannot be observed or distinguished from a team's. We take the second approach, using a unique dataset that contains the key ingredients: data on the contracts offered to each individual, *ex post* performance measures (claims by the insured), as well as transfers.

The test we propose is based on the link between insurance demand and consumption of health care. Adverse selection is present if consumers have private or unobservable information about their health state. The unobserved information links insurance choices and health care expenditures, as consumers more likely to need health care purchase more generous insurance coverage. Intuitively, the test is based on whether the link between insurance choice and health care consumption can be attributed to observables. If observables account for the link, then we can rule out the importance of unobservables in the joint insurance/health care decision. A maintained assumption of this test is that employment choice is unrelated to health status. If this is violated, the choice set of each individual may be endogenous, leading to the potential biases discussed in Section 7.

We model the unobservables by assuming that individuals receive a private signal about their future health state before making their insurance choice. After making that choice, nature determines individuals' health states. Individuals then demand health care depending on their health state, demographics, and insurance coverage. We model the discrete/continuous choice of insurance and health care consumption in the style of Dubin and McFadden (1984).

We use the single population from the 1987 National Medical Expenditure Survey (NMES), which provides relatively accurate consumption data as well as information on all policies offered to each employee, including those not chosen. We therefore have the key ingredients in this contractual relationship: policies offered, the policy chosen, and health care consumption.

We find that the link between the health insurance choice and health care consumption is mostly explained by observables. That is, we cannot reject the hypothesis that the error terms (in the two stages) are uncorrelated. There is no evidence of unobservables linking insurance status with health care demand.

The health care market provides an opportunity to compare the estimates from a discrete-choice model with experimental evidence on health care demand. We use estimates from the RAND Health Insurance Experiment (RHIE) as a basis for comparison and validation of our results. This comparison is important, since strong structural assumptions are imposed in estimating discrete-choice models, and generally there is no way of validating their implications. Our estimates of health care demand elasticities are in the range of the RHIE estimates.

Sections 2 and 3 present descriptive statistics and survey the related literature. Sections 4 and 5 present the model and discuss the estimation. Sections 6 and 7 analyze our estimates and their implications.

2. Descriptive statistics

■ We use the National Medical Expenditure Survey (NMES) collected by the Agency of Health Care Policy Research (AHCPR) in 1987. The NMES requested detailed information on health care consumption and demographics for a sample of over 13,000 households. The strength of the survey is the reliability of the health expenditure and insurance data. In general, obtaining this information is problematic, as individuals tend to forget details or fail to keep records about choices and expenses. In contrast, the NMES requested the identity of employers and insurers. Employers and insurers provided information on the individual's insurance plans, and health care providers reported actual health care consumption.

Another important feature of the data is that the NMES Health Insurance Plan Survey (HIPS) contains information on those employer-provided policies that were not chosen. This is essential information, since for estimation we want to know the characteristics of all options, whether chosen or not.

Tables 1 through 3 present the subsample of the NMES (U.S. Department of Health and Human Services, 1987) used in the estimation. It includes single individuals of working age (from 18 to 65 years old) who are employed.[2] Each observation has an associated weight to extrapolate from the sample to the whole population. The tables show weighted means. The last row shows the percentage weight of each column. For more details about the NMES, see Section 6.

Table 1 presents demographics and health insurance information. The sample of 826 individuals[3] is divided into insured and uninsured. Two observations are immediate. First, a substantial proportion of the sample is uninsured, and second, health care expenditure is important relative to total income, around 5%. This confirms the importance of the market and the need to understand who the uninsured are and why they are uninsured.

The average expenditure for the sample is $900. The distribution of expenditures is highly skewed, as expected. Insured individuals spent 50% more in health care than the uninsured ($1,019 versus $660). The gap may be caused by three factors: observable demographic differences, cost of health care, or adverse selection.

The first factor is evident from Table 1, which shows that the uninsured are younger and poorer. The second factor is the cost of treatment. If demand for health care is price sensitive, the uninsured will demand less care because they pay the full cost, whereas the insured on average pay only 12 cents per dollar spent. Finally, adverse selection may explain the gap in expenditure. Healthier individuals decline to buy insurance and then, on average, use less health care than those who, aware of their poorer health, decide to buy insurance.[4] Our goal is to disentangle the relative importance of these factors.

Since health insurance policies, which are mainly employer-provided, vary widely across individuals, we cannot conclude from Table 1 that insurance demand is income- and age-sensitive. For instance, if income and fringe benefits (including health insurance) increase with seniority, the fact that older and richer individuals are more likely to be insured may just reflect access to cheaper insurance. Thus, we have to control for the choice set each individual faces in order to estimate the determinants of the insurance status.

The difference in average premium offered to the insured and to the uninsured shows how different the choice sets are; the latter is four times higher. Some employers do not offer insurance at all, in which case the employee has to purchase insurance in the market, which is more expensive than employer-provided insurance.

Table 2 separates those offered insurance by their employer from those not offered. Smaller firms (in terms of employment) and lower-paying jobs are less likely to offer health insurance. Table 3 shows insurance status for those offered and not offered. It is clear that employer-provided insurance is cheaper, on average, than individual insurance offered in the market. As a consequence, over 90% of those offered insurance by their employer buy insurance, while over 90% of

[2] We restrict attention to single individuals; families face extremely complex choice sets, including the options offered to each spouse and coverage under different policies. We excluded people above age 65 and below age 18 because the former qualify for Medicare and the latter are likely to belong to a bigger household. Results are restricted to the sample analyzed.

[3] There are 1,008 observations that satisfy the criteria above; 182 were dropped because they showed inconsistencies, such as uninsured individuals who declined free health insurance.

[4] The self-reported health states at the bottom of Table 1 seem to contradict the presence of adverse selection, since we see that the uninsured do not report better health. But this information is hard to interpret. Insured individuals may report better health because they are getting better or more health care.

TABLE 1 **Weighted Means**

	All	Insured	Uninsured
Age	34.0	35.3	31.5
% female	44.7	50.4	33.3
Income	18,280	22,059	10,632
Total health care expenditure	901	1,019	660
Coinsurance rate (offered)	.12	.12	.13
Premium (offered)	360.1	170.6	745.2
Deductible (offered)	140	124	173
Total employees	638.53	878.01	153.38
Northeast (%)	20.40	23.09	14.93
Midwest	25.36	27.77	20.48
West	21.08	20.86	21.53
South	33.16	28.28	43.06
Hispanic (%)	7.35	4.88	12.38
Black (%)	12.05	9.28	17.68
Urban core (%)	32.27	33.37	30.04
Urban metropolitan area (%)	51.07	52.85	47.48
Nonurban (%)	16.66	13.78	22.45
Self-reported health state			
Excellent	34.3	38.0	26.8
Good	54.0	52.2	57.6
Fair	10.9	8.8	15.0
Poor	.8	.9	.6
Number of observations	826	516	310
Weight	100	67	33

Note: These are weighted means for the single population sample used in the estimation. Total health care expenditure includes the total cost of care no matter who paid for it and excludes insurance premia. Coinsurance, Premium, and Deductible are the mean of those offered to the individual. Health states are self-reported assessments that fall into the four categories above.

those not offered remain uninsured.[5] The percentages of those insured and those offered insurance vary considerably across industries and occupational categories, suggesting that the choice set is to a large extent determined by employment conditions.

In summary, there are several symptoms consistent with adverse selection: there is a large proportion of uninsured individuals, individual health insurance is substantially more expensive

[5] Adverse selection is a possible explanation for the higher insurance cost of small employers. Morrisey (1992) discusses several other reasons, such as lower incentives to subsidize and higher administrative costs. On average, insured individuals in our sample pay only 50 cents per dollar of care. If firm size is associated with higher pay (as in our sample), then larger firms have a greater incentive to subsidize insurance owing to tax advantages.

TABLE 2 **Means Conditional on Offered**

	Offered	Not Offered
Expenditure	990.5	718.4
Age	35.0	32.1
Income	21,774	11,216
Premium	152.2	—
Employer size	933.4	41.7
N	525	301
Weight	66.9	33.1

Note: Observations separated according to offered/not offered insurance by the employer. Employer size is the number of employees working for the employer.

than insurance provided by large employers, and insured individuals spend 50% more on health care than uninsured individuals. The difference in health care expenditures between the insured and uninsured could be attributed to demographics and price sensitivity. We estimate a model with the purpose of distinguishing these effects. If the estimated price and income elasticities as well as other demographics can explain the expenditure gap, then the role of adverse selection, even if present, can be judged economically insignificant.

The nature of the test is similar to the one in Cameron et al. (1988) (described in the next section) and Chiappori and Salanié (2000b). Chiappori and Salanié test the hypothesis that more comprehensive auto insurance coverage is chosen by agents with higher accident probability. They find no evidence of asymmetric information in the French car insurance market.

3. Related literature

■ **Adverse selection.** The evidence on adverse selection in health insurance markets is scarce and contradictory. Phelps (1976) found no systematic relation between predicted illness of individuals and insurance choice. Wolfe and Goddeeris (1991) estimate a health care utilization equation using a longitudinal sample of individuals more than 65 years old purchasing supplementary insurance (supplementing Medicare) in the Medigap market. They find evidence of

TABLE 3 **Means Conditional on Offered and Insured**

	Offered		Not Offered	
	Insured	Uninsured	Insured	Uninsured
Expenditure	1,005	716	1,270	654
Age	35.2	29.5	35.9	31.7
Income	22,302	11,889	17,557	10,488
Premium	138.9	400.4	757.9	—
Employer size	924.9	1,093.3	11.2	45.2
N	492	33	24	277
Weight	65.5	3.4	3.4	29.6

Note: Similar to Table 2, but with observations further separated by insured and uninsured.

adverse selection in the Medigap market, but they judge it economically nonsignificant. Marquis and Phelps (1987) find adverse selection using data from a questionnaire regarding the hypothetical purchase of supplemental insurance. Cameron et al. (1988) find adverse selection in an Australian sample. They test for adverse selection by estimating health care demand and then performing a Hausman specification test for the endogeneity of the insurance status. They find statistical dependence between the error terms in the insurance and the health care demand equations. This correlation suggests adverse selection or at least the presence of variables unobservable to the researcher that determine both insurance choice and health care consumption. Dowd et al. (1991), using an approach similar to Cameron et al., find no evidence of adverse selection in their Minnesota sample. In the market for young individuals and in the market for the elderly the evidence is mixed. Morrisey (1992) and Pauly (1986) survey the literature in detail.

□ **Health care demand.** The vast literature that estimates demand for health care uses many different approaches. The main distinction is between experimental and nonexperimental methods. The nonexperimental ones use variation in insurance plan characteristics to explain health care expenditure. The main problem is the endogeneity of health insurance that biases the estimates. Newhouse, Phelps, and Marquis (1980) provide a summary of the statistical problems present in this line of research. The most reliable evidence comes from the RAND Health Insurance Experiment. It was conducted by RAND under federal grant in the mid-1970s over six cities. Participating families (about 6,000 individuals) were randomly assigned to one of fourteen fee-for-services plans with the purpose of studying health care demand sensitivity to prices, income, deductible, and so on (Manning et al., 1987). We use these findings as benchmark comparisons for our health care demand estimates.

□ **Health insurance demand.** Researchers have used different approaches to estimate demand for health insurance. As Morrisey (p. 37, 1992) puts it: "The data have not always been well suited to the task. . . The picture of price responsiveness is not as clear for health insurance as it was for health services." Since there is no study of health insurance demand comparable to the RAND experiment, the estimates are more controversial. Morrisey (1992) surveys the early studies. Later studies such as Holmer (1984), Feldman et al. (1989), Cameron et al. (1988), and Steinberg Schone, Selden, and Zabinski (1995) estimate discrete-choice models of demand. Holmer's sample includes federal government employees, while Feldman et al. deals with 17 Minneapolis firms. In both cases all employees are insured, hence it is difficult to infer much about the behavior of the uninsured or the adverse-selection problems leading to underinsurance. Cameron et al. use Australian data to study the binary choice of purchasing supplementary insurance to the universal free basic coverage. These articles motivate the problem using a model of expected utility maximizing consumers, but they use a reduced-form model for estimation.

The structural approach we take has several advantages. It leads to interpretable parameters, some of which cannot be estimated in a reduced form. For example, it is not possible to estimate risk preferences without specifying a parametric utility function. Moreover, we obtain a direct measure of the extent of asymmetric information. The structural approach, on the other hand, makes several simplifications to handle a complex problem. Biases may arise from misspecification of key elements of the problem.

4. The model

■ The problem we have in mind is a two-stage decision made by an expected-utility-maximizing consumer. In the first stage, the consumer chooses the insurance policy that yields the highest expected utility. In the second stage, after the uncertain health state is realized, she chooses health care consumption.[7]

[7] There are some dynamic issues we are abstracting from. See Gilleskie (1998) for a structural dynamic estimation that studies the relation between absenteeism and medical care. The issues involved allow her to use data on the timing of expenditures through the sample year. In our case we have only one observation on insurance choice, so we use total expenditure over the year.

In the context of insurance, uncertainty is the essence of the problem. We model uncertainty by assuming preferences to be given by $U(m_i, h_i)$, where m_i is a composite good and h_i is health consumption. The latter is given by $h_i = x_i + s_i$, where s_i is a random health state and x_i is health care. That is, health care is a perfect substitute for the health state. The health state, s_i, is uncertain at the moment of purchasing insurance, and it can be compensated for with doctor visits, surgery, etc. Basically, s_i is a random component that affects the marginal utility from care.

Each insurance policy is associated with its main characteristics, Z_j, which include a premium, deductibles, coinsurance rates, etc. Pairs of deductibles and coinsurance rates determine the budget set faced by consumers. Thus, buying a policy means choosing a budget set for the second stage.

We capture the self-selection into policies by assuming that each consumer i gets a signal, ω_i, of her future health state, s_i, before purchasing insurance. Given their signals, consumers assess their future health status. The conditional distribution (on the signal) of the uncertain health state is used to compute expected utility under the different insurance policies. We assume that insurance choice sets are exogenous.

We now describe consumers' behavior in reverse order, starting with their health care consumption given their insurance coverage and the health state realization. We then describe optimal insurance choice conditional only on the signal but assuming optimal consumption behavior in the second stage.

□ **Second stage.** In the second stage, given a realized health state s_i and policy j, each individual i maximizes her utility by choosing the health expenditure level x_i:

$$\begin{aligned} U^*_{ij}(s_i) = U^*(y_i, s_i, Z_j) &= \max_{x_i} U(m_i, h_i) \\ &\text{subject to } m_i + C_j(x_i) = y_i - p_j, \end{aligned} \tag{1}$$

where

U^*_{ij} = indirect utility of individual i holding policy j,
y_i = income of i,
p_j = premium of policy j,
$C_j(x)$ = out-of-pocket expenditures under policy j,
Z_j = $[p_j, C_j]$ = policy j's characteristics.

Notice the difference between x_i and $C_j(x_i)$. The former represents the level of health services while the latter represents out-of-pocket expenditure. For example, if policy j has no deductible and coinsurance c_j, then $C_j(x_i) = c_j \cdot x_i$. In general, the payment is not linear in the services consumed. Most of the policies in the sample have a deductible which creates both a nonlinearity and a nonconvexity of the budget set. We shall explain in Section 5 how we incorporate this nonconvexity into the estimation algorithm.

□ **First stage.** Now we return to the first stage. After observing ω_i but prior to the full realization of the health state s_i, each consumer must choose one of the J_i insurance plans in her choice set or the outside option of remaining uninsured. Thus there are $J_i + 1$ alternatives in i's choice set.

We define in equation (2) the consumer's indirect expected utility given her private information. The indirect utility incorporates second-stage optimal behavior. To calculate expected utility conditional on ω_i and policy-specific random tastes a_{ij}, we need to integrate over the conditional distribution of s_i. For simplicity we assume the policy-specific random taste shocks, a_{ij}, enter additively, so they do not affect second-stage behavior. For each insurance policy j:

$$V_{ij}(\omega_i, a_{ij}) \equiv E(U^*_{ij}(s_i) \mid \omega_i) + a_{ij} = \int U^*(y_i, z_i, Z_j)\pi_i(dz_i \mid \omega_i, D_i) + a_{ij}, \tag{2}$$

where

π_i = distribution of the final health state s_i conditional on the signal, ω_i, and the demographics of i, D_i.
$V_{ij}(\omega_i, a_{ij})$ = i's expected utility from policy j, given signal ω_i and taste a_{ij}.

At the first stage the consumer chooses the policy that maximizes her expected utility given her private information:

$$\max_j \{V_{ij}(\omega_i, a_{ij})\} \qquad j = 0, \ldots, J_i.$$

Notice that the demographic characteristics, D_i, may affect the need for care. They enter into the model through the distribution of s_i, given by π_i. Income, on the other hand, enters through the budget constraint in (1).

5. Estimation

■ The model, like the one presented in Dubin and McFadden (1984), is a mixture of discrete and continuous choices. It involves a decision under uncertainty made by an expected utility-maximizing individual. The discrete dimension in our case involves potentially many choices, in contrast to their binary choice between electricity and gas. Furthermore, in our case the choice sets vary across individuals.

To estimate the model we use the moment conditions from both the discrete and continuous dimensions. Conditional on demographics D_i, the model predicts how individuals with different private information behave. By integrating over all those realizations of the signal that make a specific insurance policy optimal, we get the probability that such policy is chosen by each individual in the sample, conditional on her demographics. The probability of a specific choice minus an indicator function, which takes the value one if the choice was elected and zero otherwise, defines a moment condition. Moreover, conditional on having picked a specific policy in the first stage, the model predicts how much an individual spends on health care as a function of her health state, s_i. We construct expected expenditure for each plan according to the model, controlling for the endogeneity of the insurance policy choice. We use only those values of the signal that are consistent with that plan being optimal. Roughly speaking, first-stage behavior provides information about the unobserved value of the signal. The overall estimation is based on finding values of the parameters that make the model's predicted behavior match observed behavior as closely as possible.

Adverse selection is assessed by examining the variance of the private signal, which captures the link between the first and second stage unaccounted for by the demographic information. Clearly there could be individual information observed by insurers but unobserved by us. If so, we might falsely infer the existence of adverse selection in a market where information is in fact symmetric. As a consequence, our estimate is an upper bound on the extent of asymmetric information.

□ **Consumer preferences.** To estimate our model we must choose functional forms for $U(m_i, h_i)$ and for the distribution of s_i. The utility function must be flexible enough to fit the data and simple enough to provide an analytic solution for the demand function.[7] Cobb-Douglas utility, for example, is simple but places strong restrictions on price elasticity. Our solution is to use an approximation to a general, unknown utility function. Assume that the true form is $U(m_i, h_i)$. Then a second-order approximation is

$$U(m_i, h_i) \approx \phi_1 m_i + \phi_2 h_i + \phi_3 m_i \cdot h_i + \phi_4 m_i^2 + \phi_5 h_i^2.$$

Given the budget constraint $m_i = y_i - p_j - C_j(x_i)$ and $h_i = x_i + s_i$, U_{ij} is a function of x_i and s_i.

[7] If there is no analytic solution for the health care demanded, the computational burden increases, as we would have to search for the x_i that maximizes U_{ij} for all i, j at every step of the parameter search. This would require the nesting of minimization procedures.

Since most policies include deductibles together with coinsurance, we incorporate them into the model. Out-of-pocket expenditure in health becomes

$$C_j(x_i) = \begin{cases} x_i & \text{if } x_i \leq DED_j \\ DED_j + c_j(x_i - DED_j) & \text{if } x_i > DED_j. \end{cases}$$

The consumer pays the full costs of the services up to the deductible, DED_j, and above that amount she is responsible for only c_j of every dollar spent. Such a policy generates a piecewise linear budget set that is neither linear nor convex if $DED_j > 0$. To find the optimal health care consumption $x^*_{ij}(y_i - p_j, s_i, Z_j)$, we first solve the first-order condition with respect to x_i over both of the linear pieces of the budget constraint. Let $x_{ij}(y_i - p_j, s_i, 1)$ and $x_{ij}(y_i - p_j - DED_j, s_i, c_j)$ be the solutions of the first-order conditions for the case where the consumer pays the full cost of care (no deductible and coinsurance is one), and when she pays only c_j per dollar but has already paid the deductible, respectively. Solving the first-order condition gives

$$\begin{aligned} x_{ij}(y_i - p_j - DED_j, s_i, c_j) = {} & \frac{(\phi_2 + \phi_3 \cdot (y_i - p_j - DED_j) + 2 \cdot \phi_5 \cdot s_i)}{2 \cdot (c_{ij} \cdot \phi_3 - \phi_4 \cdot c_{ij}^2 - \phi_5)} \\ & - \frac{c_{ij}(\phi_1 + \phi_3 \cdot s_i + 2 \cdot \phi_4 \cdot (y_i - p_j - DED_j))}{2 \cdot (c_{ij} \cdot \phi_3 - \phi_4 \cdot c_{ij}^2 - \phi_5)}. \end{aligned}$$

Moreover, health care consumption must be within the budget set, so we must also consider the two corners, where consumption is zero or equal to $C_j(x_i) = y_i - p_j$ $(m_i = 0)$. To find optimal behavior $x^*_{ij}(y_i - p_j, s_i, Z_j)$, we compare utility at four points: at the interior solutions $x_{ij}(y_i - p_j, s_i, 1)$ and $x_{ij}(y_i - p_j - DED_j, s_i, c_j)$, and at the corners $x_{ij} = 0$ and $m_i = 0$.

This is how the model incorporates the censoring or nonnegativity constraint in the optimal predicted behavior. In contrast, reduced-form selection models need an additional equation that deals with the censoring of expenditures at zero (Goldman, 1995 and Maddala, 1985). In general, Tobit-type models include a selection equation (say between HMO versus non-HMO), a health care utilization equation, and a linear probability equation to eliminate the censoring bias (through a Mills ratio). These equations are linked only through the error structure but not by any underlying individual preferences. In our model it is the underlying preferences that determine the three levels of behavior, hence the link between them is embedded into the model.

Demand elasticities for interior solutions are

$$\varepsilon_{xc} = \frac{-x^*_{ij}(2\phi_3 - 4\phi_4 c_j) - (\phi_1 + \phi_3 \cdot s_i + 2\phi_4(y_i - p_j))}{2 \cdot (c_j \cdot \phi_3 - \phi_4 \cdot c_j^2 - \phi_5)} \cdot \frac{c_j}{x^*_{ij}}$$

$$\varepsilon_{xy} = \frac{\phi_3 - 2 \cdot c_j \cdot \phi_4}{2 \cdot (c_j \cdot \phi_3 - \phi_4 \cdot c_j^2 - \phi_5)} \cdot \frac{y_i}{x^*_{ij}}.$$

The price elasticity is a measure of moral hazard; its measurement was one of the main reasons for running the RHIE. Their results provide a reasonable target and an important specification test for our estimates.

To estimate preferences for risk, we transform the indirect utilities, defined in (1), by the CARA utility function: $-\exp\{-rU^*_{ij}(s_i)\}$. Since this is a monotonic transformation, it affects insurance choice behavior but not health care consumption. We estimate specifications with and without this transformation.

The distribution of s_i needs to be specified. It has been suggested that health care expenditures are well approximated by a log-normal distribution, but with a substantial mass at zero. Many consumers spend no money during the year, but there is a thick right tail. We use the following assumption:

$$s_i = -\exp\{K(D_i) + \omega_i + \varepsilon_i\}, \tag{3}$$

where $K(D_i)$ is a deterministic function of the demographics, $\omega_i \sim N(0, \sigma_\omega^2)$ is the signal, and $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$ is the remaining uncertainty realized after the policy has been chosen. By construction ω_i and ε_i are independent, since ε_i is the news about s_i not contained in ω_i, and any correlation should have been incorporated into the prediction or signal.

The mean of s_i is

$$E(s_i) = -\exp\{K(D_i) + \frac{\sigma_\omega^2 + \sigma_\varepsilon^2}{2}\}.$$

Thus the health state is minus a log-normal random variable. Demographics determine behavior through s_i, affecting both mean and variance. By altering the distribution of s_i they determine willingness to pay for insurance and health care consumption.

We assume the policy-specific shocks a_{ij} are independent (over both i and j) and identically distributed Type 1 Extreme Value (or Log Weibull) random variables. The logistic distribution has an analytic integral that simplifies the estimation. We do not need to numerically integrate with respect to the a_{ij}. The probability that individual i chooses option j conditional on the signal is given by

$$P_{ij}(\omega_i) = \frac{e^{V_{ij}(\omega_i)}}{\sum_{k=0}^{J_i} e^{V_{ik}(\omega_i)}}.$$

□ **Moment conditions.** This problem is similar to what Amemiya (1985) calls a multinomial generalization of a Type 5 Tobit. A selection equation determines which of several other variables are observed. We only observe x_{ij}^* if policy j is chosen. To estimate the parameters of the model we construct a method-of-moments estimator. For each person we observe a vector of indicators $I_{i0}, \ldots, I_{iJ}$, where $I_{ij} = 1$ if policy j is chosen, zero otherwise. We also observe the products of these indicators with actual expenditures: $I_{i0} \cdot x_{i0}, \ldots, I_{iJ} \cdot x_{iJ}$, where x_{ij} is actual expenditure under policy j. Using predicted and actual values, we form a vector of prediction errors:

$$u_i(\theta, D_i) = \begin{pmatrix} P_{i0}(\theta, D_i) \cdot x_{i0}^e(\theta, D_i) - I_{i0} \cdot x_{i0} \\ \vdots \\ P_{iJ}(\theta, D_i) \cdot x_{iJ}^e(\theta, D_i) - I_{iJ} \cdot x_{iJ} \\ P_{i1}(\theta, D_i) - I_{i1} \\ \vdots \\ P_{iJ}(\theta, D_i) - I_{iJ} \end{pmatrix},$$

where θ is the vector of parameters we want to estimate, and

$$x_{ij}^e(\theta, D_i) = \int \int_{\{\omega | \{V_{ij}(\omega_i, a_{ij}) \geq V_{ik}(\omega_i, a_{ik})\ \forall k\}} x_{ij}^*(y_i - p_j, K(D_i) + \omega + \varepsilon, Z_j) f(\omega) f(\epsilon) d\omega d\epsilon$$

is the predicted expenditure of individual i under policy j. If the model is correctly specified, the error terms have mean zero at the true parameter vector, θ_0. The full set of moment conditions is obtained by assuming that there exists a set of instruments W_i such that $E(u_i \mid W_i, \theta_0) = 0$. The instruments are discussed in the next section. Then

$$G(\theta_0) = E(W \otimes u(\theta_0, D_i) \mid W_i, D_i) = 0.$$

Given a parameter vector θ, which includes the parameters of the distribution of s_i as well as those of the utility function, the estimation goes as follows: For each individual i we predict her optimal expenditures under each plan. Then we integrate with respect to ε_i to get the expected utility under each plan j conditional on the private signal ω_i, $V_{ij}(\omega_i, a_{ij})$. We then find P_{ij} by integrating over values of ω_i and a_{ij} that make j the optimal choice. The second-stage consumption prediction x_{ij}^e is obtained by integrating x_{ij}^* (with respect to ω_i) over the region where the j option is optimal.[8]

Finally, in most discrete-choice model applications the choice set is the same for every individual; that is, the same set of cars or computers is available to every individual in the population. This is not the case here, where choice sets depend on employment-provided insurance, and are indexed by i. This does not create any econometric problems, but it does necessitate careful treatment of the moment conditions.

It is worth mentioning that variability in choice sets helps the estimation, as it creates variation of plan characteristics across otherwise similar individuals. Table 3 hints at how demand elasticities can be estimated. Each person evaluates the expected utility of each insurance option, including no insurance. Because similar consumers face different premia, they will make different insurance choices and will end up facing different coinsurance rates. Heterogeneity[9] in the choice set across individuals provides the variability to identify consumer preferences.

The moment conditions cannot just enumerate the different options and add all the error terms of each option, since options across consumers have nothing in common. We construct the moment conditions by type of insurance in the following way. We first calculate choice probabilities and expected expenditure for each plan. Then we sum choice probabilities and expenditure for three segments: no insurance, fee-for-service (FFS), and HMO. The probability that a person chooses an HMO, for example, is taken to be the sum over the probabilities of individual HMO plans in the person's choice set. Expenditures are calculated using a probability weighted average. The prediction error $u_i(\theta, D_i)$ consists of five elements as defined above.

6. Data

■ We use a subsample from the NMES that includes single individuals of working age (from 18 to 65 years old) who are employed. As a consequence, our findings on the presence of adverse selection are restricted to the population studied. Restricting the sample to a subset of individuals involves the risk of making the results nonrepresentative of the population as a whole, and it has the drawback that interesting questions about household behavior are left unanswered.

The number of policies offered to the 826 individuals in the sample range from zero to eight. For estimation we reduced the number of available options. Each person's choice set contains the chosen plan, the outside good, and at most two HMO plans and two non-HMO or FFS plans. Thus each person has at most four options plus no insurance. Nonchosen alternatives were randomly selected from the reported options of each person. Very few cases were affected by this aggregation, as few consumers are offered more than three policies. There are 301 individuals who were not offered employer-provided insurance, while 273 were offered a single policy, 100 two policies, 83 individuals three policies, and 69 four or more policies. Those not offered plans by their employers either purchase a private plan or remain uninsured. We do not observe the choice sets of those individuals who are not offered insurance and remain uninsured. For each of these individuals we randomly assigned an actual plan chosen by those who did buy a private plan. Sample statistics were shown in Table 1. Coinsurance is calculated by averaging across many covered services, when they differed. We use six instruments: age, age squared, sex, race dummies, region dummies, and a constant.[10]

---

[8] We use Monte Carlo integration, with at least 25 points in the support of each dimension. There are only two dimensions of integration, ε and ω.

[9] Choice sets are endogenous if firms offer different combinations of insurance and wages. We assume employment is exogenous, since we have no data to endogenize job selection. We find no evidence of firms specializing in employing different types of workers. We discuss endogeneity in Section 7.

[10] Notice that all these variables are part of the observable information about each insured and are included among the demographics in estimation. If these variables are correlated with unobservables, they will pick up those effects, but we are interested in their "reduced-form" role. Namely, we want the observables to contain as much information as possible about an individual's type.

TABLE 4 Main Estimates

	Parameter	Estimate (1)	t-statistic	Estimate (2)	t-statistic	Estimate (3)	t-statistic
1	σ_ω (signal)	.52	2.50	.12	.43	.06	.11
2	σ_ε (noise)	.99	2.87	1.75	7.00	1.95	4.22
3	ϕ_1 (M)	1	—	1	—	1	—
4	ϕ_2 (H)	−.94	−.22	−.70	−.23	−.61	−.20
5	ϕ_3 (M·H)	.04	.42	.11	1.40	.36	.45
6	ϕ_4 (M^2)	−.001	−.70	−.015	−13.12	−.004	−.17
7	ϕ_4 (H^2)	−7.27	−.64	−6.39	−1.13	−4.70	−.44
8	Age			−.05	−1.82	−.02	−1.07
9	Age^2			.001	3.24	.001	2.74
10	Female			.67	1.95	1.15	2.53
11	Northwestern			−.17	−.56	−.03	−.61
12	Nonmetro			−.57	−1.33	−.56	−1.22
13	Black			1.28	2.28	1.36	2.83
14	Clerical			−.23	−.32	.04	.92
15	Constant	−1.07	−1.33	−2.37	−3.26	−3.97	−4.95
16	HMO coin	2.55	1.61	2.66	3.15	2.24	3.83
17	$r(-e^{-rV})$					.04	.93
$J \sim \chi^2$		52.16		29.79		43.78	

Note: The first two rows present the estimated standard errors of the signal and of the noise around the signal. The next five rows present the coefficients of the quadratic utility function (with the linear term in M normalized to one). Rows 8 to 13 present the demographics in the function $K(D_i)$, i.e., the deterministic component of the health state. Row 15 presents the estimated HMOs "fictitious" coinsurance. Row 16 presents the CARA transformation parameter. t-statistics are based on a covariance matrix of the estimators computed using numerical gradients for a 10% increase in each parameter.

7. Results

■ Preferences are described by the coefficients of the quadratic utility function, the effect of the demographics on the distribution of the health state, ($K(D_i)$ (equation 3)), and the coefficient of risk aversion. $K(D_i)$ is assumed to be a linear function of age, age squared, sex, region of the country (dummies for Northwestern and nonmetropolitan area), a race dummy (for black), and a dummy for white-collar worker. We included the demographic information that appeared relevant in reduced-form regressions of health care consumption.

Results are shown in Table 4. The estimates include both variances, σ_ω and σ_ε (the first two rows), the utility function parameters, ϕ_i, the effect of the demographics in the distribution of s_i (reported in rows 8 to 15), the HMO dummy (row 16), and the risk-aversion coefficient, r.

We normalize the coefficient on m in the utility function to be one. This normalization is inconsequential, since Bernoulli utility functions are equivalent up to affine transformations. The first column presents basic estimates without controlling for demographics. The second set of estimates adds demographics. The third specification adds the CARA transformation.

The J test of the overidentifying restrictions is equal to the optimized objective function times the sample size (826). The test statistic is distributed approximately χ^2 with degrees of freedom equal to the number of overidentifying restrictions. There are 30 moment conditions arising from six instruments. The three specifications in Table 4 have, respectively, 8, 15, and 16 parameters, yielding 22, 15, and 14 overidentifying restrictions. The value of the test statistic for the model with demographics (column 2) is close to the 5% critical value of 25.0, but it falls in the rejection region. In terms of fit, the model without demographics is much poorer than the one with demographics, as expected. The model with the CARA transformation also does worse (the models are not nested). For the rest of this section we shall refer to the model in column 2. We discuss the risk aversion estimates in Section 7.

The sign of ϕ_2 seems peculiar. The reason ϕ_2 is negative is that the health state s_i, and thus h_i, are negative. But it is easy to verify that the utility function is concave and increasing (in the relevant range) in both arguments. For a high realization of s_i, utility is satiated in h_i, capturing the fact that some people do not spend any money on health care.

We calculate the effect of the demographics on health care consumption through their effect on the distribution of s_i. Adding five years of age implies 10% higher expenditure, or about \$75. Women are predicted to spend about \$500 more, while in the sample women on average spent \$450 more than men. Blacks are predicted to spend an extra \$1,300, while in the sample they spend \$1,050 more. Note that the model's predictions, unlike sample averages, are calculated holding other variables (income, insurance, and other demographics) constant.

HMOs provide managed care. They deserve special treatment, as consumers do not freely make their choices based on coinsurance and deductibles. The way we treat them is by replacing the coinsurance data in the demand function for HMOs with an estimated parameter. Intuitively, it is a constant that reflects a coinsurance level that fits the consumption patterns of HMO enrollees. We have in mind that HMOs have a pool of resources to treat their patients and allocate these resources according to their needs or health draws. The parameter measures the shadow price of the resource constraint consistent with observed expenditures. The implementation assumes that consumption, x_i, is determined by the HMO according to that shadow cost, but HMO enrollees get the treatment for free. Thus, behavior is determined by the HMO according to the shadow cost of care, while utility depends on the level of care and the actual out-of-pocket expenditure. Estimates of this parameter show that HMOs succeed in reducing demand. The parameters suggest that consumers act as though coinsurance rates were over 100%. The estimated coinsurance is large, but demand elasticity (reported in the next section) is low, so that HMO patients get about 10% less health care overall. This parameter, however, might be picking up other differences between HMOs and other plans.[11]

□ **Model performance.** To evaluate the fit of the model, we compare predicted and observed behavior in Tables 5 and 6. In Table 5 we define the model's predicted insurance plan as the one with the highest probability of being chosen. The i, j cell gives the number of instances in which plan i was predicted and plan j chosen. Diagonal elements are correct predictions. Row totals are aggregate predictions. Column totals are aggregate observed choices. The model closely predicts aggregate shares of each. At the individual level it correctly predicts about 66% of the cases.

Table 6 shows that the model underpredicts mean health care expenditure by about 20%. While 52% of the sample spend less than \$200, the model predicts that proportion to be 47%. Figures 1 and 2 show predicted and observed log expenditures. Figure 2 is a simulated sample using individual characteristics and model estimates. The spikes at zero show the frequency of zero expenditure. While the distributions are similar, the model predicts thinner tails.[12] In particular, the model underpredicts the probability of zero consumption.

Health care demand elasticities are reported in Table 6. The elasticities in the column labelled

[11] A possible concern is that HMOs attract healthier consumers. To test this possibility, we used the fact that 25% of the individuals offered HMO were not offered non-HMO. We reestimated the model allowing for an additional dummy for those offered only HMOs. The coefficient on the dummy is not significant and has the wrong sign, showing no selectivity bias (which is consistent with Dowd et al. (1991)).

[12] Right tail: about 1.6% of the sample spend more than \$10,000. Their mean expenditure is \$22,231 and the maximum expenditure is \$64,000. The model predicts 1% of the sample spends more than \$10,000, with a mean expenditure of \$17,076 and a maximum of \$52,840.

TABLE 5 **Insurance Choice Model Fit**

	Observed			
	Uninsured	Non-HMO	HMO	Total
Predicted				
Uninsured	203	92	12	307
Non-HMO	100	292	38	430
HMO	7	29	53	89
Total	310	413	103	826
Correct predictions: 66.3%				

Note: Cells present the number of observation that fall into each specific predicted/observed pair. Observations in the diagonal are those with insurance status correctly predicted by the model.

"actual" are from the RHIE. Estimated price elasticity (our measure of moral hazard) is $-.18$ (evaluated at no insurance). The numbers are small but similar to the RHIE findings, which range from $-.1$ to $-.2$ (Manning et al. (1987)). Estimated income elasticity is .51, higher than the experimental findings.

Preferences toward risk are determined by both r and the concavity of U. In spite of the insignificant estimate of r we think it is of interest to analyze the point estimate. The estimates imply a coefficient of relative risk aversion of 1.8. Zeldes' (1989) estimation of intertemporal consumption reports a utility function implying a coefficient of relative risk aversion of 2.3, close to the numbers we find.

□ **Asymmetric information.** The main purpose of estimation is to study the importance of asymmetric information. We want to test whether unobservables link the two stages. Those unobservables, potentially private information, are captured by σ_ω, the standard deviation of the signal. A large σ_ω means that consumers receive dispersed signals. For a given level of demographics, those receiving a low draw of ω_i (healthier) are willing to pay less for insurance, since they expect their health to be relatively good. In turn they will end up spending less on health care. In contrast, those with high draws (less healthy) are willing to pay more for insurance and in turn end up on average with a higher marginal utility from health care. If σ_ω is low or zero, then consumers receive basically the same information, and no link, besides that captured by observable demographics, is present between demand for insurance and demand for health care. If demographics explain most of the link, then unobservables play no important role, as the risk type can be categorized by demographics.

There are two identification concerns. First, we need to explain how adverse selection, captured by σ_ω, is separated from moral hazard. The fact that more generous insurance leads to

TABLE 6 **Expenditures Goodness of Fit**

	Predicted	Actual
Mean expenditure	749	980
Insured	901	1,112
Uninsured	496	661
Price elasticity	.18	".2"
Income elasticity	.51	".2"

Note: Comparisons of model predictions with sample values. "Actual" elasticities are those from the RHIE.

FIGURE 1

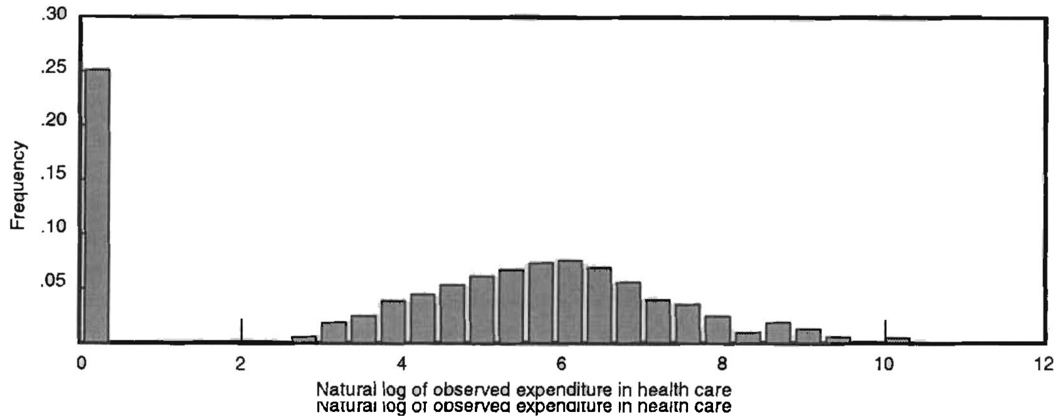

higher demand for health care could also be attributed to price sensitivity, i.e., moral hazard. The key to identification is the heterogeneity of choice sets across individuals. If all individuals in the sample faced the same insurance policies, we would not be able to tell the two apart, for both explanations would be observationally equivalent. We would just observe two classes of people. The higher expenditures of the insured could be attributed to lower coinsurance or to adverse selection. In contrast, since in our data choice sets differ across individuals, we observe similar individuals facing similar policies (in terms of coinsurance and deductibles) with different premia. Then similar consumers end up facing different coinsurance rates. Price sensitivity is identified using coinsurance variability across individuals. On the other hand, if similar consumers facing similar choice sets make different decisions, then we attribute the difference to unobservables or to a dispersed signal. Another way to put it is that coinsurance is an endogenous variable, whose coefficient in the health care utilization equation would be picking up both the moral hazard and the adverse-selection effects. But since different individuals face different choice sets, the premium at which insurance is offered to them becomes an instrument for the coinsurance. Premium offered, taken as exogenous, will be correlated with coinsurance but uncorrelated with the individual health state.

The argument that the price (moral hazard) effect is nonparametrically identified can be made more formally. Suppose for simplicity that the choice is binary (purchasing insurance or remaining uninsured) and that offered policies differ only in their premium and coinsurance. We can define a cutoff type $\underline{\omega}(p_i, c_i)$ such that all $\omega \geq \underline{\omega}(p_i, c_i)$ prefer to buy insurance over remaining uninsured. Since $\underline{\omega}(p_i, c_i)$ increases in both arguments (for more expensive insurance the indifferent consumer must be less healthy), we can find all the pairs (p_i, c_i) such that $\underline{\omega}$ is constant. These negatively related p_i and c_i describe the different policies that attract consumers with the same health type. Notice that as we move on the (p_i, c_i) locus, only c_i affects the second-stage behavior (p_i is irrelevant at that stage). Moreover, the effect of c_i on consumption is not contaminated by endogeneity, since all individuals come from the same distribution (as they have the same signal or type).

The second identification concern is how σ_ω is distinguished from σ_ε. The variables ε and ω play identical roles in s (see (3)). However, σ_ω and σ_ε can be separately identified by their distinct roles linking the insurance and health care stages. We have explained in the previous paragraph how a large σ_ω makes those with low draws behave differently from those with high draws, creating a positive correlation between willingness to pay for insurance and health care expenditure across the sample. In contrast, a high σ_ε affects all the sample uniformly, increasing each person's willingness to pay for insurance (as the uncertainty in health care need increases) and at the same time increasing their mean expenditure.[13]

[13] An increase in σ_ε increases mean expenditure because of the censoring of health care demand at zero. σ_ε affects the right tail of expenditures but not the left.

FIGURE 2

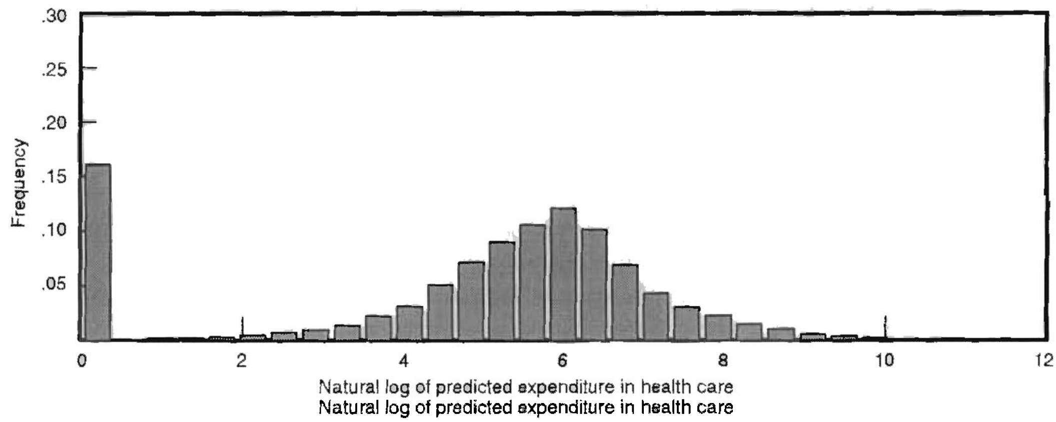

According to the estimates in columns 2 and 3 of Table 4 (including demographic information), we cannot reject the hypothesis that σ_ω is zero.[14] There is no evidence of superior information by buyers. In other words, the link between insurance choices and expenditures is accounted for by observables.[15]

The next question is whether the model fails to show asymmetric information because there is none or because the model is inappropriate. We claim above on theoretical grounds that σ_ω is identified. A practical test of the power of the model in identifying unobservable factors is the following exercise. We found that some demographics such as age and race are significant determinants of health care consumption. We can omit them from the estimation, treating them as unobservables. We do not use all the relevant individual information, and we expect the model to reflect this through σ_ω. Column 1 shows estimates without controlling for demographics. The estimate of σ_ω increases by a factor of four, and we reject equality to zero at the 1% level. This suggests that the model is able to estimate the asymmetry in information.

□ **Robustness.** To test robustness we estimated two alternative models. First, we estimated a simple alternative model that resembles Cameron et al. (1988). We estimate a health care utilization equation with demographics and health insurance status as explanatory variables. Then we test the independence between insurance status and the error term.

The simplified model is appropriate for the application in Cameron et al., where all individuals in their Australian sample faced the same binary choice. It does not account for the main features of our sample, such as the more complex, individual-specific choice sets. The model neglects the censoring of expenditures at zero, so the estimates are likely to be biased. Incorporating the additional features of the problem would bring us back to our proposed model. The purpose is to check the robustness of results in a simplified model.

We perform a Hausman specification test comparing the OLS estimates to 2SLS estimates. For the 2SLS estimation we instrument for insurance status (insured or uninsured) with insurance choice set characteristics (number of policies offered, premium, etc.) These are valid instruments under the assumption that employment decisions, and therefore the characteristics of offered plans, are exogenous.

The null hypothesis is that the error term is uncorrelated with health insurance status. Under the null, OLS estimates are consistent and efficient, but inconsistent under the alternative. 2SLS estimates are consistent under both hypotheses but not efficient under the null. The test is based

[14] The testing is complicated by the parameter at the boundary value. But since we are testing only one parameter, then the χ^2 tests applicable to parameters in a corner collapse to the usual t-tests (Wolak (1989)).

[15] It can be argued that there might be adverse selection since race is correlated with expenditure while insurers cannot discriminate based on race. But our purpose is to test for adverse selection as an intrinsic market characteristic. We are not concerned with whether current laws lead to adverse selection, or whether firms do not use all available information.

on

$$m = (\beta_{2SLS} - \beta_{OLS})'(V(\beta_{2SLS}) - V(\beta_{OLS}))^{-1}(\beta_{2SLS} - \beta_{OLS}),$$

where $V(\cdot)$ represents an estimated covariance matrix. Under the null, m is distributed $\chi^2(k)$, where k is the number of parameters (in this case 10). The parameters hardly differ, and $m = .43$ while the critical value at the 5% confidence level is 18.3. In concurrence with our previous finding, the null is not rejected under this simple model.

A key element of the model is the distribution of the signal. As a second robustness test, we replaced the log normal-like distribution with a Bernoulli random variable with support zero and a, where the outcome a arises with probability p. The discrete signal might better capture a situation in which most of the population have the same information while a small proportion generate large expenditures. The point estimates are reasonable, implying an 8% probability of spending an extra \$360. The conclusions are unchanged in that the estimates a and p are not significantly different from zero.

□ **Potential biases.** The main potential source of bias is the endogeneity of employment, and hence of the choice set. If firms specialize in different compensation packages, workers are likely to self-select into firms based on both their observables and unobservables. The latter is the source of bias.

To assess the importance of choice set endogeneity, we performed probit regressions of the probability that the individual's employer offered insurance. We want to test the extent of self-selection into jobs based on observables, which will give us an indication of the extent of selection based on unobservables. Table 7 contains selected results for three specifications. The first includes twelve industry and ten occupation dummy variables along with demographics and income. The second drops the industry and occupation variables, and the third also drops income variables. We do not report coefficients on the industry and occupation dummies.

In the third specification, some of the demographic variables, age and race in particular, are significant. When income is added (second specification), the estimated coefficients on the demographic variables shrink, and none of them is significant. Moreover, notice that income increases the probability of being offered insurance. Under the endogeneity assumption, income should have the opposite sign, as those in greater need of health insurance would give up wages to get better insurance. In contrast, the positive relation suggests that health insurance, as a fringe benefit, is highly related to monetary compensation. When industry and occupation dummies are added, the log-likelihood increases dramatically. While a couple of industries (transportation and public administration) offer better access to insurance, most of the variation seems to be across occupations.

We test the joint significance of demographic variables using a likelihood-ratio test. The test statistic is $2(\ell(\theta^U) - \ell(\theta^R)) \sim \chi^2(r)$, where ℓ denotes the log-likelihood function evaluated at unrestricted and restricted parameter values and r is the number of restrictions. The unrestricted model (1) is compared with two restricted specifications, which we do not report. The null hypothesis is that the coefficients on the omitted variables are all equal to zero. In the first we omitted eleven income, age, female, race, and health status variables. This yields a likelihood value of -312.67 and a test statistic value of 62. The critical value for a one-sided test at the 5% level is 19.68. In the second restricted specification we replace the two income variables. This yields a likelihood value of -284.49 and a test-statistic value of 5.64. The critical value for a one-sided test at the 5% level is 16.92. We reject the null when income variables are omitted, but we cannot reject the null when only the other demographics are omitted.

While nothing conclusive can be drawn from these results, they are not suggestive of endogeneity, at least based on observable factors. After we control for income and job characteristics, as proxies for skill, there appears to be little self-selection into jobs that offer insurance. Notice also that self-reported health status is not significant in any of the specifications. Under the selection assumption, either through observables or unobservables, these variables should matter.

TABLE 7 Probit Estimation of the Probability of Being Offered Insurance

Parameter	Estimate (1)	*t*-statistic	Estimate (2)	*t*-statistic	Estimate (3)	*t*-statistic
Age	.0107	1.17	.0174	1.65	.0419	4.14
Age^2	−.0001	−.93	−.0002	−1.38	−.0005	−3.47
Female	.0559	.562	.1158	1.04	.1859	1.74
Female · age	−.0007	−.26	.0012	.71	−.0012	−.402
Northeast	.0565	1.33	.0942	1.84	.1576	3.30
Midwest	.0757	1.98	.0738	1.54	.1165	2.59
West	.0333	.74	.0294	.53	.0684	1.33
Hispanic	−.0533	−.99	−.0640	−.98	−.1460	−2.33
Black	.0022	.05	−.0045	−.10	−.0531	−1.20
Core	.0152	.32	.0141	.26	.1148	2.25
Metro	.0588	1.39	.0234	.48	.1116	2.40
Excellent	−.0428	−.232	−.1012	−.56	.1092	.68
Good	−.0723	−.41	−.1512	−.86	.0281	.17
Fair	−.0589	−.30	−.1753	−.90	−.0838	−.48
Income	.0223	6.97	.0415	11.37	—	—
Income^2	−.0003	−5.94	−.0004	−8.31	—	—
Total employees	.0003	5.46	—	—	—	—
Log-likelihood/Observations	−281.67/822		−414.16/826		−498.06/826	

If selection based on unobservables is present, the nature and direction of possible biases are not transparent, due to the nonlinearity of the model. We might expect the endogeneity to cause an upward bias in the moral hazard estimates and a downward bias in the income elasticity. Individuals with worse health draws are attracted by firms offering more generous insurance coverage. Our price elasticity estimate is low at .18. So if a bias is present, it is probably not large.

The main parameter of interest, σ_ω, is likely to be contaminated since it is jointly identified with the moral hazard estimate. There is no obvious way in which the parameter should be affected. We expect upward bias in the moral hazard estimates to bias σ_ω downward, as the two are competing explanations for why the insured spend more in health care than the uninsured in the second stage.

8. Conclusions

■ This article estimates a discrete/continuous model that integrates both health insurance and health care demand, taking into account buyers' private information. Our main goal is to study the scope of adverse selection in health insurance markets. We do not find evidence of asymmetric information. The link between insurance choice and health care demand can be explained mostly by observables, leaving a small and statistically insignificant role for unobservables.

Having failed to find significant evidence of adverse selection, we need an explanation for the gap in health care expenditure between the insured and the uninsured, and for the fact that employer-provided insurance is substantially cheaper, especially for larger employers. According to the estimates, the gap in expenditure is due to observable demographic differences and to price sensitivity. There may be several factors that make insurance provided by large employers cheaper. For instance, fixed administrative costs and bargaining power may give larger firms a cost advantage. We also find in Section 2 that larger employers in our sample offer a greater health insurance subsidy. Larger firms tend to pay higher wages, and on average their employees face higher marginal tax rates, giving a greater incentive to subsidize.

An alternative contract-theoretic explanation, which deserves more empirical attention, is based on Cochrane (1995). In his model of health insurance, buyers cannot commit to long-term contracts. Lack of long-term contracts can explain why small firms and individual buyers cannot find affordable insurance. Absent long-term commitment, healthy individuals (and small employers) would drop coverage, leaving only bad draws in the pool. In contrast, large employers, who base their decisions on the average draw of all their employees, are less likely to withdraw from the pool. Evidence on the importance of commitment in long-term contracts is reported in Hendel and Lizzeri's (2000) study of the life insurance industry. This alternative explanation has received very little empirical attention (Crocker and Moran (1998)).

References

AKERLOF, G.A. "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, Vol. 84 (1970), pp. 488–500.

AMEMIYA, T. *Advanced Econometrics.* Cambridge, Mass.: Harvard University Press, 1985.

CAMERON, A.C., TRIVEDI, P.K., MILNE, F., AND J. PIGGOTT. "A Microeconometric Model of the Demand for Health Care and Health Insurance in Australia." *Review of Economic Studies*, Vol. 55 (1988), pp. 85–106.

CHIAPPORI, P.A. AND SALANIÉ, B. "Testing Contract Theory: A Survey of Some Recent Work." Mimeo, Department of Economics, University of Chicago, 2000a

—— AND ——. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, Vol. 108 (2000b), pp. 56–78.

COCHRANE, J.H. "Time-Consistent Health Insurance." *Journal of Political Economy*, Vol. 103 (1995), pp. 445–473.

CROCKER, K. AND MORAN, J. "Contracting with Limited Commitment: Evidence from Employment Based Insurance Contracts." Mimeo, University of Michigan, 1998.

DOWD, B., FELDMAN, R., CASSOU, S., AND FINCH, M. "Health Plan Choice and the Utilization of Health Care Services." *Review of Economics and Statistics*, Vol. 73 (1991), pp. 85–93.

DUBIN, J.A. AND MCFADDEN, D.L. "An Econometric Analysis of Residential Electric Appliance Holding and Consumption." *Econometrica*, Vol. 52 (1984), pp. 345–362.

FELDMAN, R., FINCH, M., DOWD, B., AND CASSOU, S. "The Demand for Employment-Based Health Insurance Plans." *Journal of Human Resources*, Vol. 24 (1989), pp. 115–142.

GILLESKIE, D.B. "A Dynamic Stochastic Model of Medical Care Use and Work Absence." *Econometrica*, Vol. 66 (1998), pp. 1–45.

GOLDMAN, D. "Managed Care as a Public Cost-Containment Mechanism." *RAND Journal of Economics*, Vol. 26 (1995), pp. 277–295.

HENDEL, I. AND LIZZERI, A. "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance." NBER Working Paper no. W7470, 2000.

HOLMER, M. "Tax Policy and the Demand for Health Insurance." *Journal of Health Economics*, Vol. 3 (1984), pp. 203–221.

MADDALA, G.S. "A Survey of the Literature on Selectivity Bias as it Pertains to Health Care Markets." In L.F. Rossiter and R.M. Scheffler, eds., *Advances in Health Economics and Health Services Research.* Greenwich, Conn.: JAI Press, 1985.

MANNING, W.G., NEWHOUSE, J.P., DUAN, N., KEELER, E.B., LEIBOWITZ, A., AND MARQUIS, M.S. "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment." *American Economic Review*, Vol. 77 (1987), pp. 251–277.

MARQUIS, M.S. AND PHELPS, C.E. "Price Elasticity and Adverse Selection in the Demand for Supplementary Health Insurance." *Economic Inquiry*, Vol. 25 (1987), pp. 299–313.

MORRISEY, M. *Price Sensitivity in Health Care: Implications for Health Care Policy.* Washington, D.C.: The National Federation of Independent Business Foundation, 1992.

NEWHOUSE, J., PHELPS, C., AND MARQUIS, S. "On Having Your Cake and Eating It Too: The Econometric Problems of Estimating Health Services." *Journal of Econometrics*, Vol. 13 (1980), pp. 365–390.

PAULY, M.V. "Taxation, Health Insurance, and Market Failure in the Medical Economy." *Journal of Economic Literature*, Vol. 24 (1986), pp. 629–675.

PHELPS, C. "The Demand for Reimbursement Insurance." In R.N. Rossett, ed., *The Role of Health Insurance in Health Services Sector.* New York: National Bureau of Economic Research, 1976.

ROTHSCHILD, M. AND STIGLITZ, J.E. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, Vol. 90 (1976), pp. 630–649.

STEINBERG SCHONE, B., SELDEN, T., AND ZABINSKI, D. "Estimating a Structural Model of Insurance Choice and Medical Care Utilization." Working Paper, AHCPR, 1995.

U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES, AGENCY FOR HEALTH CARE POLICY AND RESEARCH. *National Medical Expenditure Survey*. Washington, D.C.: U.S. GPO, 1987.

WOLAK, F.A. "Testing Inequality Constraints in Linear Econometric Models" *Journal of Econometrics*, 41 (1989), pp. 203–235.

WOLFE, J.R. AND GODDEERIS, J.H. "Adverse Selection, Moral Hazard, and Wealth Effects in the Medigap Insurance Market." *Journal of Health Economics*, Vol. 10 (1991), pp. 433–459.

ZELDES, S.P. "Consumption and Liquidity Constraints: An Empirical Investigation" *Journal of Political Economy*, Vol. 97 (1989), pp. 305-346.

LINKED CITATIONS

You have printed the following article:

Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey
James H. Cardon; Igal Hendel
The RAND Journal of Economics, Vol. 32, No. 3. (Autumn, 2001), pp. 408-427.
Stable URL:
http://links.jstor.org/sici?sici=0741-6261%28200123%2932%3A3%3C408%3AAIIHIE%3E2.0.CO%3B2-N

This article references the following linked citations. If you are trying to access articles from an off-campus location, you may be required to first logon via your library web site to access JSTOR. Please visit your library's website or contact a librarian to learn about options for remote access to JSTOR.

[Footnotes]

[1] **Taxation, Health Insurance, and Market Failure in the Medical Economy**
Mark V. Pauly
Journal of Economic Literature, Vol. 24, No. 2. (Jun., 1986), pp. 629-675.
Stable URL:
http://links.jstor.org/sici?sici=0022-0515%28198606%2924%3A2%3C629%3ATHIAMF%3E2.0.CO%3B2-4

[7] **A Dynamic Stochastic Model of Medical Care Use and Work Absence**
Donna B. Gilleskie
Econometrica, Vol. 66, No. 1. (Jan., 1998), pp. 1-45.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28199801%2966%3A1%3C1%3AADSMOM%3E2.0.CO%3B2-A

[11] **Health Plan Choice and the Utilization of Health Care Services**
Bryan Dowd; Roger Feldman; Steven Cassou; Michael Finch
The Review of Economics and Statistics, Vol. 73, No. 1. (Feb., 1991), pp. 85-93.
Stable URL:
http://links.jstor.org/sici?sici=0034-6535%28199102%2973%3A1%3C85%3AHPCATU%3E2.0.CO%3B2-2

References

NOTE: *The reference numbering from the original has been maintained in this citation list.*

LINKED CITATIONS

The Market for "Lemons": Quality Uncertainty and the Market Mechanism
George A. Akerlof
The Quarterly Journal of Economics, Vol. 84, No. 3. (Aug., 1970), pp. 488-500.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28197008%2984%3A3%3C488%3ATMF%22QU%3E2.0.CO%3B2-6

A Microeconometric Model of the Demand for Health Care and Health Insurance in Australia
A. C. Cameron; P. K. Trivedi; Frank Milne; J. Piggott
The Review of Economic Studies, Vol. 55, No. 1. (Jan., 1988), pp. 85-106.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28198801%2955%3A1%3C85%3AAMMOTD%3E2.0.CO%3B2-5

Testing for Asymmetric Information in Insurance Markets
Pierre-André Chiappori; Bernard Salanié
The Journal of Political Economy, Vol. 108, No. 1. (Feb., 2000), pp. 56-78.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28200002%29108%3A1%3C56%3ATFAIII%3E2.0.CO%3B2-J

Time-Consistent Health Insurance
John H. Cochrane
The Journal of Political Economy, Vol. 103, No. 3. (Jun., 1995), pp. 445-473.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28199506%29103%3A3%3C445%3ATHI%3E2.0.CO%3B2-6

Health Plan Choice and the Utilization of Health Care Services
Bryan Dowd; Roger Feldman; Steven Cassou; Michael Finch
The Review of Economics and Statistics, Vol. 73, No. 1. (Feb., 1991), pp. 85-93.
Stable URL:
http://links.jstor.org/sici?sici=0034-6535%28199102%2973%3A1%3C85%3AHPCATU%3E2.0.CO%3B2-2

An Econometric Analysis of Residential Electric Appliance Holdings and Consumption
Jeffrey A. Dubin; Daniel L. McFadden
Econometrica, Vol. 52, No. 2. (Mar., 1984), pp. 345-362.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198403%2952%3A2%3C345%3AAEAORE%3E2.0.CO%3B2-Y

LINKED CITATIONS

The Demand for Employment-Based Health Insurance Plans
Roger Feldman; Michael Finch; Bryan Dowd; Steven Cassou
The Journal of Human Resources, Vol. 24, No. 1. (Winter, 1989), pp. 115-142.
Stable URL:
http://links.jstor.org/sici?sici=0022-166X%28198924%2924%3A1%3C115%3ATDFEHI%3E2.0.CO%3B2-1

A Dynamic Stochastic Model of Medical Care Use and Work Absence
Donna B. Gilleskie
Econometrica, Vol. 66, No. 1. (Jan., 1998), pp. 1-45.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28199801%2966%3A1%3C1%3AADSMOM%3E2.0.CO%3B2-A

Managed Care as a Public Cost-Containment Mechanism
Dana P. Goldman
The RAND Journal of Economics, Vol. 26, No. 2. (Summer, 1995), pp. 277-295.
Stable URL:
http://links.jstor.org/sici?sici=0741-6261%28199522%2926%3A2%3C277%3AMCAAPC%3E2.0.CO%3B2-W

Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment
Willard G. Manning; Joseph P. Newhouse; Naihua Duan; Emmett B. Keeler; Arleen Leibowitz
The American Economic Review, Vol. 77, No. 3. (Jun., 1987), pp. 251-277.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28198706%2977%3A3%3C251%3AHIATDF%3E2.0.CO%3B2-S

Taxation, Health Insurance, and Market Failure in the Medical Economy
Mark V. Pauly
Journal of Economic Literature, Vol. 24, No. 2. (Jun., 1986), pp. 629-675.
Stable URL:
http://links.jstor.org/sici?sici=0022-0515%28198606%2924%3A2%3C629%3ATHIAMF%3E2.0.CO%3B2-4

Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information
Michael Rothschild; Joseph Stiglitz
The Quarterly Journal of Economics, Vol. 90, No. 4. (Nov., 1976), pp. 629-649.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28197611%2990%3A4%3C629%3AEICIMA%3E2.0.CO%3B2-N

Consumption and Liquidity Constraints: An Empirical Investigation
Stephen P. Zeldes
The Journal of Political Economy, Vol. 97, No. 2. (Apr., 1989), pp. 305-346.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28198904%2997%3A2%3C305%3ACALCAE%3E2.0.CO%3B2-M